



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/782,314	02/19/2004	Michael R. Smith	CIS0211US	7633
33031 7590 02/15/2011 CAMPBELL STEPHENSON LLP 11401 CENTURY OAKS TERRACE BLDG. H, SUITE 250 AUSTIN, TX 78758				
EXAMINER JAKOVAC, RYAN J				
ART UNIT		PAPER NUMBER		
2445				
MAIL DATE		DELIVERY MODE		
02/15/2011		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/782,314

Filing Date: February 19, 2004

Appellant(s): SMITH ET AL.

Shawn Doman

For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 11/22/2010 appealing from the Office action mailed 05/21/2010.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is incorrect. The status in the brief states that no claims have been cancelled, however, claims 4, 18, and 28-37 have been cancelled.

(4) Status of Amendments After Final

The Appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The Appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

The following is a listing of the prior art of record relied upon in the rejection of claims under appeal:

1. US 2001/0014097 to Beck, et al (“Beck”).
2. TCP/IP Illustrated, Volume 1: The Protocols (“TCP/IP”).
3. US 2005/0083933 to Fine, et al (“Fine”).
4. US 6,735,205 to Mankude, et al (“Mankude”).

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

1. **Claims 1-3, 5-7 are rejected** under 35 U.S.C. 103(a) as being unpatentable over U.S. 2001/0014097 to Beck et al (hereinafter Beck) in view of TCP/IP Illustrated, Volume 1: The Protocols (hereinafter TCP/IP).

Regarding independent claim 1, Beck teaches a system comprising:

a virtual link bundle comprising a plurality of communication links [Beck, fig. 2, 7. See also, the abstract, and at least paragraphs (0003-0004), (0008), (0011), (0027), and (0034), links to multiple nodes managed using a cluster address.], wherein

the plurality of communication links is configured to couple a virtual network device to a first network device external to the virtual network device [Beck, fig. 2, fig. 7, cluster (i.e. virtual network device) is coupled to external router (i.e. first network device).];

a first end of each of the communication links is configured to be coupled to a first network device [Beck, The communications links including the links between the processor nodes of the cluster and the router are shown in Fig. 2.];

a second end a first one of the communication links is configured to be coupled to a first virtual network device sub-unit within a virtual network device [Beck, Fig. 2 discloses Processor node B (i.e. first virtual network device sub-unit) which is a node inside the cluster (i.e. virtual network device). The communications links are show connected via the subnet s1 between the network router and the processor nodes.];

a second end of a second one of the communication links is configured to be coupled to a second virtual network device sub-unit within the virtual network device [Beck, Fig. 2 discloses Processor node C (i.e. second virtual network device sub-unit) which is a node inside the cluster (i.e. virtual network device). The communications links are shown connected via the subnet s1 between the network router and the processor nodes.];

the first one of the communication links and the second one of the communications links provide redundant connections between the first network device and the first virtual network device [Beck, Fig. 2, multiple communications links between cluster and processor nodes. See also, [0026], multiple subnets couple the processing nodes. Further, providing redundant connections between networked elements is a well known technique in the art used to provide a reliable network. Providing redundant connections between networked elements would have been obvious to one of ordinary skill in the art at the time of the invention.];

the first network device comprises a plurality of ports [Beck, fig. 2, router comprises a plurality of ports.];

each of the ports is configured to communicate packets with a respective client [Beck, fig. 2, router communicating packets.].

Beck does not explicitly disclose the first network device is configured to append a header to a packet before sending the packet to the virtual network device via one of the communication links. However, the network router of Beck (i.e. the first network device) sends packets to via one of the communication links of fig. 2 and discloses the use of TCP/IP network protocols [Beck, (0007).].

TCP/IP discloses that a header includes identification one of the ports having received the packet [TCP/IP, pg. 2-4, the format of the TCP header includes a source port number (i.e. identifies the receiving port).].

Therefore it would have been obvious to one of ordinary skill in the art at the time of invention to for Beck's router to append a header to a packet before sending the packet to the virtual network device via one of the communication links where the header identifies a one of the ports having received the packet in order to identify the source port [TCP/IP, pg. 2-4.]. It would have been further obvious since including this information in a header comprises the use of standard TCP/IP protocol which is a well known standard in the art.

Regarding dependent claim 2, the combination of Beck and TCP/IP teaches the system of claim 1, further comprising the first network device, wherein

the first network device is configured to select a communication link of the plurality of communication links on which to send a particular packet [Beck, (0064-0074). See (0009), A data packet is sent to the network for delivery to a particular node (i.e. along a particular link to the node).].

Regarding dependent claim 3, The combination of Beck and TCP/IP teaches the system of claim 2, wherein

each packet sent between the virtual network device and the first network device is sent via only a one of the communication links [Beck, Paragraph (0030), A destination address is

specified for a particular node, thus packets are sent via the specific link associated with an individual processor.].

Regarding dependent claim 5, The combination of Beck and TCP/IP teaches the system of claim 1, further comprising the first virtual network device sub-unit, wherein

the first virtual network device sub-unit is configured to identify whether a one of the communication links is coupled to another virtual network device sub-unit within the virtual network device [Beck, Paragraph (0011), When a link has failed, the address of the processor node is acquired by another processor node for the duration of the failure.].

Regarding dependent claim 6, The combination of Beck and TCP/IP teaches the system of claim 1, further comprising the first virtual network device sub-unit [Beck, Fig. 2, Processor node B] and the second virtual network device sub-unit [Beck, Fig. 2, processor node C], wherein

the first virtual network device sub-unit and the second virtual network device sub-unit are configured to communicate packets with each other via a virtual network device link [Beck, Paragraph (0027), Processor nodes distribute packets within the cluster.].

Regarding dependent claim 7, The combination of Beck and TCP/IP teaches the system of claim 1, wherein

the communication links are configured to be managed as a single link [Beck, Paragraph (0004), Cluster alias addresses are used to make the cluster appear to be a single node. See (0027), (0034).].

2. **Claims 8-17, 19-20, 22-27, 38, 40-48, 50-58, 60-67 are rejected** under 35 U.S.C. 103(a) as being unpatentable over U.S. 2001/0014097 to Beck et al (hereinafter Beck) in view of US 2005/0083933 to Fine et al (hereinafter Fine).

Regarding independent claim 8, Beck teaches a system comprising:

a first virtual network device sub-unit [Beck, Fig. 2, Processor Node .] comprising:

a first interface [Beck, Fig. 2, interface of processor node A]; and

a controller coupled to the first interface and configured to forward packets received via the first interface [Beck, paragraph (0027), the processor node distributes packets to other nodes within the cluster.], wherein

the first interface is identified by a first logical identifier [Beck, fig. 2, processor node A is identified by the cluster alias address. See paragraphs (0004), (0027), (0034).],

a second interface is identified by the first logical identifier [Beck, fig. 2, processor node B is identified by the cluster alias address. See paragraphs (0004), (0027), (0034).],

an interface bundle comprises the first interface and the second interface [Beck, Fig. 2, grouping of processor nodes called a cluster with interconnected communication links.), and

the second interface is comprised in a second virtual network device sub-unit [Beck, fig. 2, interface of processor node B.],

the controller is configured to detect whether a packet was received via a virtual network device link [Beck, fig. 4, packet is received and a determination is made whether it was sent to the cluster (i.e. via the virtual network device links of fig. 2.). See also (0034), (0027).],

a first end of the virtual network device link is configured to be coupled to the first virtual network device sub-unit [Beck, fig. 2.],

a second end of the virtual network device link is configured to be coupled to the second virtual network device sub-unit [Beck, fig. 2.].

Beck does not expressly disclose the first interface is configured to filter out the packet from a packet flow being sent via the first interface if the packet was received via the virtual network device link.

However, Fine from the same field of endeavor teaches the packet is filtered to prevent a client in a multi-access network from receiving duplicate packets [Fine, paragraph (0005).].

Therefore, it would have been obvious to one skilled in the art at the time of the invention was made to prevent packets from being resent as taught by Fine into one of the processor nodes as taught by Beck. The motivation for doing this is to improve the network efficiency by not overloading the system.

Regarding dependent claim 9, the combination of Beck and Fine teaches the system of claim 8, further comprising the second virtual network device sub-unit [Beck, Fig. 2 discloses a system including multiple processor nodes (i.e. second virtual network device sub-unit.)].

Regarding dependent claim 10, the combination of Beck and Fine teaches the system of claim 9, wherein

the first virtual network device sub-unit is configured to maintain consistent forwarding information with the second virtual network device sub-unit [Beck, paragraph (0027), the processor node distributes packets to other nodes within the cluster.].

Regarding dependent claim 11, the combination of Beck and Fine teaches the system of claim 10, wherein

the controller is configured to perform control protocol processing for the first interface according to a routing protocol running on the interface bundle [Beck, paragraph (0076), Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster.],

the controller is configured to provide information generated when performing the control protocol processing to a secondary controller comprised in the second virtual network device sub-unit [Beck, paragraph (0053-0054), The Mbuf chain data structure is sent across a cluster interconnect to a processor node.], and

the secondary controller is configured to use the information to manage the second interface [Beck, paragraph (0055), The Mbuf chain is queued for service of packets on the second node.].

Regarding dependent claim 12, the combination of Beck and Fine teaches the system of claim 8, wherein

the controller is configured to lookup a destination address of a first packet in a lookup table [Beck, paragraph (0040-0041), The receiver listens on a destination port number. The processor node looks up receiver applications in a lookup table.], and

if the lookup table returns the first logical identifier, the first virtual network device sub-unit is configured to prioritize sending the first packet via the first interface over sending the first packet via the second interface [Beck, Paragraph (0041-0044), a database is maintained from which a node is selected based on selection weights which include IP address (i.e. logical identifier).].

Regarding dependent claim 13, the combination of Beck and Fine teaches the system of claim 12, wherein

if the lookup table returns the first logical identifier, the first virtual network device sub-unit is configured to send the first packet via the first interface instead of sending the packet via the second interface, unless one or more of the first interface and a link coupled to the first interface are failed [Beck, paragraph (0076), Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.].

Regarding dependent claim 14, the combination of Beck and Fine teaches the system of claim 13, wherein

the first virtual network device sub-unit comprises a plurality of interfaces, more than one of the interfaces are each comprised in the interface bundle, and the more than one of the interfaces comprises the first interface [Beck, paragraph (0062), each node may contain more than one interface. Fig. 7 discloses nodes having multiple interfaces. Nodes with multiple interfaces are comprised in the cluster (i.e. interface bundle).].

Regarding dependent claim 15, the combination of Beck and Fine teaches the system of claim 14, wherein

if each respective communication link coupled to the more than one of the interfaces fails, the first virtual network device sub-unit is configured to forward the first packet via the second interface comprised in the second virtual network device sub-unit [Beck, paragraph (0076), Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.].

Regarding dependent claim 16, The combination of Beck and Fine teaches the system of claim 8, wherein

the first virtual network device sub-unit [Beck, Fig. 2, Processor Node A] is coupled to the second virtual network device sub-unit by a virtual network device link [Beck, Fig. 2 discloses the network links between the processors in the cluster including the network subnet.].

Regarding dependent claim 17, the combination of Beck and Fine teaches the system of claim 16, wherein

the first virtual network device sub-unit is configured to learn that a source address of the second packet is behind the first interface, in response to receiving a second packet via the virtual network device link [Beck, paragraph (0009), The packet is sent from one node to another over the network.].

Regarding independent claim 19, the combination of Beck and Fine teaches a system comprising:

a virtual link bundle [Beck, Fig. 2 discloses subnet S1 connected to processor nodes of a cluster. The subnet and the connections to the processor nodes comprise the virtual link bundle.];

a first virtual network device sub-unit [Beck, Fig. 2, Processor node B], wherein

the first virtual network device sub-unit is configured to detect whether a packet was received via a virtual network device link [Beck, fig. 4, packet is received and a determination is made whether it was sent to the cluster (i.e. via the virtual network device links of fig. 2.)],

a first end of the virtual network device link is configured to be coupled to the first virtual network device sub-unit [Beck, fig. 2],

a second end of the virtual network device link is configured to be coupled to the second virtual network device sub-unit [Beck, fig. 2], and

a second virtual network device sub-unit [Beck, Fig. 2, Processor node C], wherein

a first interface of the first virtual network device sub-unit is coupled to the virtual link bundle [Beck, Fig. 2, Processor node B with processor interface 20b and associated links. See also fig. 7.],

a second interface of the second virtual network device sub-unit is coupled to the virtual link bundle [Beck, Fig. 2, Processor node C with processor interface 20C and associated links. See also fig. 7.], and

each of the first interface and the second interface is identified by a first logical identifier [Beck, (0004), (0027), (0034), Fig. 2, interfaces of cluster nodes are identified by the cluster alias address.].

Beck does not expressly disclose the first interface is configured to filter out the packet from a packet flow being sent via the first interface if the second packet was received via the virtual network device link.

However, Fine from the same field of endeavor teaches the packet is filtered to prevent a client in a multi-access network from receiving duplicate packets [Fine, paragraph 0005].

Therefore, it would have been obvious to one skilled in the art at the time of the invention was made to prevent packets from being resent as taught by Fine into one of the processor nodes as taught by Beck. The motivation for doing this is to improve the network efficiency by not overloading the system.

Regarding dependent claim 20, the combination of Beck and Fine teaches the system of claim 19, further comprising:

a network device coupled to the first virtual network device sub-unit and the second virtual network device sub-unit by the virtual link bundle [Beck, Fig. 2 discloses a node (i.e. network device) connected to the processor nodes (i.e. first and second virtual network devices).].

Regarding dependent claim 22, the combination of Beck and Fine teaches the system of claim 19, wherein

a primary controller comprised in the first virtual network device sub-unit is configured to perform control protocol processing for the first interface according to a routing protocol running on the virtual link bundle [Beck, Paragraph [0007], Nodes distribute TCP packets to other processing nodes for servicing]),

the primary controller is configured to send information generated by performing the control protocol processing to a secondary controller comprised in the second virtual network device sub-unit, and the secondary controller is configured to use the information to manage the second interface [Beck, Paragraph (0055), The Mbuf chain is sent to a node from a second node and queued for service on the second node.].

Regarding dependent claim 23, the combination of Beck and Fine teaches the system of claim 19, wherein

the first virtual network device sub-unit is configured to lookup a destination address of a packet in a lookup table [Beck, paragraph (0041), the node accesses a lookup table in accordance with which port and a node (i.e. destination address) is listening on for incoming packets.], and

if the lookup table returns the first logical identifier, the first virtual network device sub-unit is configured to prioritize sending the packet via the first interface over sending the packet via the second interface [Beck, paragraph (0076), Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.].

Regarding dependent claim 24, the combination of Beck and Fine teaches the system of claim 23, wherein

each of a plurality of interfaces comprised in the first virtual network device sub-unit is coupled to a respective communication link comprised in the virtual link bundle, and the interfaces comprise the first interface [Beck, paragraph (0062), each node may contain more

than one interface. Fig. 7 discloses nodes having multiple interfaces. Nodes with multiple interfaces are comprised in the cluster (i.e. interface bundle).].

Regarding dependent claim 25, the combination of Beck and Fine teaches the system of claim 24, wherein

if each respective communication link coupled to the interfaces fails, the first virtual network device sub-unit is configured to send the packet via the second interface comprised in the second virtual network device sub-unit [Beck, paragraph (0076), Nodes in the cluster establish databases containing the network layer addresses used by each processor node in the cluster. If the processor node crashes, another node takes over.].

Regarding dependent claim 26, the combination of Beck and Fine teaches the system of claim 23, wherein

the first virtual network device sub-unit [Beck, Fig. 2, Processor Node A.] is coupled to the second virtual network device sub-unit by a virtual network device link [Beck, Fig. 2 discloses the network links between the processors in the cluster including the network subnet.].

Regarding independent claims 38, 48, and 58, the combination of Beck and Fine teaches a method, system, and a computer readable medium comprising:

sending a first packet via a first link of a virtual link bundle if a destination identifier associated with the first packet is associated with the virtual link bundle [Beck, fig. 2 displays links between a plurality of processors. The processors communicate over the communication links disclosed in fig. 2. See also (0039-0041).]; and

sending a second packet via a second link of the virtual link bundle if a destination identifier associated with the second packet is associated with the virtual link bundle [Beck, paragraph (0009), When a receiving node determines which processor node to send to, it sends the data packet over the network for delivery to the processor node. See also fig. 2.], wherein

the destination identifier associated with the first packet identifies a destination [Beck, (0063-0071), destination information is included with the packets.],

the destination identifier associated with the second packet identifies the destination [Beck, (0063-0071), destination information is included with the packets.],

the first link is coupled to a first virtual network device sub-unit (Beck, Fig. 2 discloses a cluster containing processor node B (first virtual network sub-unit) which is connected via links associated with the subnet depicted in Fig. 2.), and

the second link is coupled to a second virtual network device sub-unit (Beck, Fig. 2 discloses a cluster containing processor node C (second virtual network device sub-unit) which are connected via links associated with the subnet depicted in Fig. 2.).

Regarding dependent claims 40, 50, and 60, the combination of Beck and Fine teaches the method of claim 39, the system of claim 49, and the computer readable medium of claim 59 further comprising:

appending a header to the first packet [Beck, paragraph (0009), *The data packet's header* is modified before it is delivered across the network to the processor node (i.e. via the first link).],

wherein the header identifies which port of a plurality of ports received the first packet [Beck, paragraph (0030), The TCP/IP header identifies the source node, the destination node, the sending port, the destination port, and the protocol being used.], and

the sending the first packet via the first link comprises sending the header via the first link [Beck, (0009, 0040).].

Regarding independent claims 41, 51, and 61, the combination of Beck and Fine teaches a method, system, and the computer readable medium comprising:

receiving a packet [Beck,(0027), (0034), (0039), fig. 4.], wherein

a destination identifier for the packet identifies an interface bundle [Beck, paragraph (0004), Cluster alias addresses are used to make the cluster appear to be a single node.], and

the interface bundle comprises a first interface [Beck, Fig. 2 discloses a cluster comprising processor nodes, each with a processor interface. The processor interface of processor node B is the first interface.];

detecting whether the packet was received via a virtual network device link [Beck, fig. 4, packet is received and a determination is made whether it was sent to the cluster (i.e. via the virtual network device links of fig. 2.)]; wherein

a first end of the virtual network device link is coupled to a first virtual network device sub-unit [Beck, fig. 2.],

a second end of the virtual network device link is coupled to a second virtual network device sub-unit [Beck, fig. 2.];

Beck does not expressly disclose and filtering out the packet from a packet flow being sent via the first interface if the packet was received via the virtual network device link.

However, Fine from the same field of endeavor teaches the packet is filtered to prevent a client in a multi-access network from receiving duplicate packets [Fine, paragraph 0005.].

Therefore, it would have been obvious to one skilled in the art at the time of the invention was made to prevent packets from being resent as taught by Fine into one of the processor nodes as taught by Beck. The motivation for doing this is to improve the network efficiency by not overloading the system.

Regarding dependent claims 42, 52, and 62, the combination of Beck and Fine teaches the method of claim 41, the system of claim 51, and the computer readable medium of claim 61 further comprising:

sending the packet via the first interface if the packet was not received via the virtual network device link [Beck, Paragraph (0058), Nodes forward packets to each other via the cluster interconnect.].

Regarding dependent claims 43, 53, and 63, the combination of Beck and Fine teaches the method of claim 42, the system of claim 52, and the computer readable medium of claim 62 further comprising:

maintaining consistency between a lookup table comprised in the first virtual network device sub-unit and a second lookup table comprised in the second virtual network device sub-unit [Beck, paragraph (0076), each processor node within the cluster establishes a database containing the network layer address used by each of the processor nodes in the cluster.].

Regarding dependent claims 44, 54, and 64, the combination of Beck and Fine teaches the method of claim 42, the system of claim 52, and the computer readable medium of claim 62 further comprising

performing control protocol processing for the interface bundle at a primary controller comprised in a first virtual network device sub-unit [Beck, Paragraph [0007], Nodes distribute TCP packets to other processing nodes for servicing.],

wherein the first interface is comprised in the first virtual network device sub-unit [Beck, Fig. 2 discloses a cluster comprised of nodes (i.e. first virtual network device sub-unit) each of which has a processor interface.].

Regarding dependent claims 45, 55, and 65, the combination of Beck and Fine teaches the method of claim 44, the system of claim 54, and the computer readable medium of claim 64 further comprising:

managing a second interface of the second virtual network device sub-unit in response to information generated by the performing the control protocol processing [Beck, paragraph [0055], The Mbuf chain is queued for service of packets on the second node.], wherein

the second interface is comprised in the interface bundle [Beck, Fig. 2 discloses a cluster (i.e. interface bundle) comprising processor node C with processor interface 20c.].

Regarding dependent claims 46, 56, and 66, the combination of Beck and Fine teaches the method of claim 45, the system of claim 55, and the computer readable medium of claim 65 further comprising:

looking up a destination address of a second packet in a lookup table [Beck, paragraph (0030-0036), The TCP/IP header identifies the source node, the destination node, the sending port, the destination port, and the protocol being used. TCP port numbers are used to designate queues into which arriving packets are placed for service by nodes. Paragraph (0076), each processor node has a database containing the network layer addresses used by each processor in order to take over in case of a failure.], and

if the lookup table returns the destination identifier, sending the sending packet via the first interface of the first virtual network device sub-unit instead of sending the packet via the second interface of the second virtual network device sub-unit [Beck, paragraph (0076), if a processor node crashes, another takes over for it.].

Regarding dependent claims 47, the combination of Beck and Fine teaches the method of claim 41, further comprising:

learning that a source address of the packet is behind a local interface, in response to receiving the packet via the virtual network device link [Beck, paragraph(0063), nodes on the same subnet communicate directly. Paragraph (0066) discloses that one node of sends another a message. The receiving node subsequently responds to the sender which indicates that the receiving node is configured to learn that the source address is the sending node. See also (0009), See also [0039-0041].].

Regarding dependent claim 57, the combination of Beck and Fine teaches the system of claim 51, further comprising:

means for learning that a source address of the packet is behind a local interface, in response to receiving the packet via the virtual network device link [Beck, paragraph(0063), nodes on the same subnet communicate directly. Paragraph (0066) discloses that one node of sends another a message. The receiving node subsequently responds to the sender which indicates that the receiving node is configured to learn that the source address is the sending node. See also (0009), See also [0039-0041].].

Regarding dependent claim 67, the combination of Beck and Fine teaches the computer readable medium of claim 61, wherein the program instructions are further executable to:

learn that a source address of the packet is behind a local interface, in response to detecting reception of the packet via the virtual network device link [Beck, paragraph(0063), nodes on the same subnet communicate directly. Paragraph (0066) discloses that one node of sends another a message. The receiving node subsequently responds to the sender which

indicates that the receiving node is configured to learn that the source address is the sending node. See also (0009), See also [0039-0041].].

3. **Claims 21, 39, 49, and 59 are rejected** under 35 U.S.C. 103(a) as being unpatentable over the combination of Beck and Fine in view of U.S. 6,735,205 to Mankude et al (hereinafter Mankude).

Regarding claim 21, the combination of Beck and Fine teaches the system of claim 20, teaches wherein

the network device is configured to send a packet via the selected one of the communication links [Beck, paragraph (0027),nNodes distribute packets to other nodes within the cluster.].

Beck and Fine do not expressly disclose the network device is configured to use a hash-based load-sharing algorithm to select one of a plurality of communication links comprised in the virtual link bundle.

However, Mankude discloses the network device is configured to use a hash-based load-sharing algorithm to select one of a plurality of communication links comprised in the virtual link bundle [Mankude, Col. 7, line 10-25, In order to select a server node to forward the packet to, the system hashes the source address of the client. The hashing selects an entry to identify a server node within a clustered computing system.].

Therefore it would have been obvious to one of ordinary skill in the art at the time of invention to combine teaches wherein the network device is configured to use a hash-based load-sharing algorithm to select one of a plurality of communication links comprised in the virtual link bundle as taught by Mankude with the system of the combination of Beck and Fine in order to select a server node to forward a packet to [Mankude, Col. 7, line 10-25].

Regarding claims 39, 49, and 59, the combination of Beck and Fine teaches the method of claim 38, the system of claim 48, and the computer readable medium of claim 58 further comprising:

selecting the first link from a plurality of links comprised in the virtual link bundle [Beck, Paragraph [0009], When a receiving node determines which processor node to send to, it sends the data packet over the network for delivery to the processor node.],

Beck and Fine do not expressly disclose wherein the selecting comprises performing a hash-based algorithm.

However, Mankude discloses wherein the selecting comprises performing a hash-based algorithm (Mankude, Col. 7, line 10-25, In order to select a server node to forward the packet to, the system hashes the source address of the client. The hashing selects an entry to identify a server node within a clustered computing system.).

Therefore it would have been obvious to one of ordinary skill in the art at the time of invention to combine Mankude with Beck and Fine in order to select a server node to forward a packet to [Mankude, Col. 7, line 10-25]).

(10) Response to Argument

In the Argument, Appellant argued in substance that:

(A) That the term “virtual link bundle” is a term of art having a specific meaning to those of skill in the art. Appellant argues that the cited references fail to disclose a virtual link bundle.

As to point (A), the Examiner respectfully disagrees. There is no evidence of record to suggest that the Appellant is correct in asserting that “virtual link bundle” is a term of art. The Appellant has not clearly set forth a definition of the term in the specification. Although the Appellant’s specification provides examples for the term, it fails to provide an explicit definition. There is no evidence of record establishing the term as a term of art used in this particular way. Appellant has failed to meet the requirement to act as his or her own lexicographer with regards to the term “virtual link bundle”. There is no evidence or record besides the Appellant’s assertion that “virtual link bundle” is a specific term of art.

All the limitations of the Appellant’s claims are met. In particular, Beck discloses a virtual link bundle since Beck discloses links coupling a cluster to a router which are managed as a single logical link. Beck’s cluster comprises a plurality of nodes and is analogous to the Appellant’s virtual network device. Beck’s router is analogous to the Appellant’s network device. Packets are directed to the cluster via a “cluster alias address [Beck, paragraph (0027), (0034), (0004)]. See for example the Appellant’s specification in paragraph (0036) describing a virtual link bundle as the links coupling a virtual network device to a network device which can be operated as a single logical link. Therefore Beck discloses a “virtual link bundle”.

(B) The cited references fail to disclose a network device configured to select one of a plurality of communication links on which to send a packet.

As to point (B), the Examiner respectfully disagrees. Beck discloses a network device sending packets in paragraph (0034). The packets are disclosed as being sent to the cluster as a whole via the cluster alias address as well as to an individual node. Each individual node is associated with its own address which may be on a distinct subnet [see Beck, fig. 7]. Therefore packets sent to different subnets or different interfaces are sent via a different link which is selected by the sender of the packet. Moreover, Beck discloses that the packets may be sent to multiple interfaces of the cluster explicitly [Beck, (0044)]. Each of the nodes can be used as paths (i.e. links) to the subnet [Beck, (0074)]. Each node advertises itself to the router as a network route (i.e. link) to the subnet. The routers in Beck use a subnet mask to determine the subnet portion (i.e. to select a link) for the destination of the packet [Beck, (0069)]. Moreover, Beck in paragraph (0070) discloses that the router may choose multiple paths in sending packets to its destination.

(C) The cited references fail to disclose multiple communications links configured to be managed as a single link.

As to point (C), the Examiner respectfully disagrees. Beck discloses multiple links to the individual nodes of the cluster to be managed as a single link using the cluster alias address. Packets are sent to the cluster alias address and subsequently handled by an individual node. See at least paragraph (0004) of Beck: "To make a cluster appear to be a single processor node, it should have a single network layer address. Such a network layer address is referred to as a "cluster alias address". That cluster alias address should not be tied to one specific node within the cluster but rather should be collectively associated with all the processor nodes."

(D) The cited references fail to disclose a first and second interface that are both identified by a first logical identifier.

As to point (D), the Examiner respectfully disagrees. Beck discloses that the first and second interfaces of the processor nodes A and B of the cluster are both identified by the cluster alias address. See at least paragraph (0027) and fig. 7 of Beck.

(E) The cited references fail to disclose a virtual network device sub-unit configured to learn that a source address of a packet is behind an interface on a separate virtual network device sub-unit, in response to receiving the packet via a virtual network device link.

As to point (E), the Examiner respectfully disagrees. Beck discloses nodes of the cluster sending packets to other processor nodes of the cluster. The nodes are analogous to the virtual network device sub-units. The nodes on the same subnet communicate directly [Beck, (0063)]. Paragraph (0066) discloses that one node sends another a message. The receiving node subsequently responds to the sender which indicates that the receiving node is configured to learn that the source address is the sending node (i.e. configured to learn that a source address of a packet is behind an interface on a separate virtual network device). Further, paragraph (0067) discloses that normal IP routing methods will be used to send packets where a cluster alias address is configured in a virtual subnet. Normal IP routing includes the source address which is sent to the receiver. Fig. 3 shows a TCP-IP packet issued from the cluster.

(F) The cited references fail to disclose a virtual network device sub-unit configured to prioritize sending a packet via first interface of an interface bundle over sending the packet via a second interface of the interface bundle.

As to point (F), the Examiner respectfully disagrees. Beck discloses that each node in the cluster is assigned a selection weights which indicates a processor node's capacity for servicing new connections, in relation to the other processor nodes in the cluster. The nodes use this information to prioritize which node will service the packet [Beck, (0042-0043).]. Servicing the packet comprises sending the packet via the chosen nodes respective interface. Further, Beck discloses in at least paragraph (0076) a takeover method by which one node in the cluster is prioritized to take over sending packets of another node of the cluster.

(G) The rejection of claims 21, 39, 49, and 59 under USC 103(a) as purportedly being unpatentable over Beck and Fine in view of Mankude is unfounded since these claims are patentable at least by virtue of depending from allowable base claims.

As to point (G), the Examiner disagrees for the reasons provided above.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this Examiner's answer.

For the above reasons, it is believed that the rejection should be sustained.

Respectfully submitted,

/R.J./
Examiner, Art Unit 2445
[Ryan Jakovac]

/Andrew Caldwell/
Supervisory Patent Examiner, Art Unit 2445

CONFEREES:

/Andrew Caldwell/
Supervisory Patent Examiner, Art Unit 2445

/Ajay Bhatia/
Primary Examiner, Art Unit 2445